

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE 15 Jan 1999		3. REPORT TYPE AND DATES COVERED FINAL TECHNICAL	
4. TITLE AND SUBTITLE Packaging and Cooling Considerations for Cryogenic Operation of Microprocessors				5. FUNDING NUMBERS C DABT63-97-C-0072	
6. AUTHOR(S) Jonathan P. Knauth Barry Penswick					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Infrared Components Corporation 2306 Bleecker Street, Utica, NY 13501				8. PERFORMING ORGANIZATION REPORT NUMBER JK150199-01	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) DARPA/PM/DSO Directorate of Contracting PO BOX 12748, Ft. Huachuca, AZ 85670-2748				10. SPONSORING / MONITORING AGENCY REPORT NUMBER CDRL A001	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION / AVAILABILITY STATEMENT Unrestricted				12b. DISTRIBUTION CODE 19990122 101	
13. ABSTRACT (Maximum 200 words) This report contains a compilation of requirements and challenges for the cryogenic packaging of computer processors. A brief discussion of computer processing speed enhancements from sub-ambient cooling is included for use in cost/benefit analysis. A thermal analysis is presented which estimates the required cooling capacity needed for operation of microprocessors at sub-ambient temperatures over the range of from -40°C to 77 Kelvins. An analysis of electrical interconnection heat load is presented which concludes that this is the key challenge for the efficient design of sub-ambient computer systems, since reduction in heat load must be balanced with electrical impedance requirements, and the number of interconnections is quite large. A review of refrigeration cycles is presented and specific thermodynamic cycles are selected for separate portions of the operational temperature range. As this is a report for a development contract that was canceled due to lack of available funding, summary conclusions, including estimated cost data for specific cooled computer systems, are not included.					
14. SUBJECT TERMS Cryogenic Cooling, Cryocooler, Electronics Cooling, Electrical Interconnections, Cold Computing				15. NUMBER OF PAGES 49	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL		

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

Section I – Chip Thermal, Mechanical and Electrical Packaging

1.0 Computing Performance

1.1 Expected Performance Increase with reduced temperature

The program's goal is to develop computer cooling and packaging technologies that will reduce the operating temperature of microprocessors as a means of improving computer performance. Specifically, operating speeds can be increased with decreasing temperatures.

The power consumption of a computer chip follows the general form of:

$$P \propto fV^2$$

Where f is the system clock frequency and V is the operating voltage. At lower temperatures, efficient cooling will permit greater frequencies of operation. A benefit of lower temperature operation is that the electrical resistance decreases proportionately, which acts to decrease voltage requirements slightly (perhaps as much as 10% without modification to the standard processor). The net effect is improved operational speed at lower temperature. A detailed analysis of this phenomenon and the various mechanisms and strategies for improving processor speed are beyond the scope of this effort, however an understanding of the trends is essential to the development of system thermal, electrical and cost models. A number of researchers have performed work in this area, and sample data points are presented in Figure 1a below, with a curve fit of all values shown in Figure 1b.




Figure 1a - Performance increase of CMOS with decreased temperature




Figure 1b - Composite curve fit of CMOS performance increase data

1.2 Cost competition of a cooled system with parallel processors

One upper limit on the acceptable system cost for a cooled computer system is the cost of the short-term alternative of merely adding additional processors to boost performance. The actual feasibility of this approach is somewhat dependent upon the nature or the computing task. The relationship between performance speedup with multiple processors is a function of the fraction of the problem that can be performed in a parallel fashion following Amdahl's law:

$$\text{Speedup} = 1 / (\text{serial portion} + \text{parallel portion} / \text{Number of processors})$$

$$\text{serial portion} + \text{parallel portion} = 1$$

The portion of a typical application that must be performed in a serial fashion can be implied from the rule of thumb that a second processor boosts performance by 78%, as reported by

Jerome¹. From this value, we compute the typical serial portion to be 0.12 and extrapolate the relative value of additional processors to system performance for serial fraction values up to 15% as shown in Figure 1c. This approach presupposes that the applications are not optimized for parallel processing, which is a good assumption for low orders of parallelism.




Figure 1c - Parallel processor speedup under Amdahl's Law

Therefore, the relative cost of an additional processor is:

$$\text{Real Cost} = \text{Actual Cost of additional processor implementation} \div \text{Performance Increase}$$

This formula may be used for cost vs. performance analysis for decision making of cooled processor systems vs. parallel processor systems.

1.3 Future state of computing

The Semiconductor Research Corporation has issued a report entitled the "National Technology Roadmap for Semiconductors", in which the industry consensus on the future state of computing is presented. The graphs that follow are based on the estimates from the most recent report, 1997.

The minimum logic voltage (V_{dd}), is expected to be progressively lowered with future generations of processors as shown in Figure 1e:

¹ Jerome, Marty, *PC Computing*, November 1997

Innovative Techniques for Cooling Computer Chips to Cryogenic Temperatures
Infrared Components Corporation
January 13, 1999 - DABT63-97-C-00721




Figure 1d - Anticipated Logic Voltage Level (V_{dd}) in out years




Figure 1e - Chip-to-board interconnect speeds in out years




Figure If - Number of chip-to-board interconnects




Figure Ig - Processor Heat Dissipation

2.0 Low Temperature Processor Packaging

2.1 Temperature Definitions

Considerable differences exist between the commercial market place and the scientific community on the definition of various of "low" temperature refrigeration system operating ranges. For the current effort we will emphasize the commercial definitions which are briefly defined below with an example of each noted:

Conventional Refrigeration

Covers the temperature ranges associated with air conditioning and conventional refrigeration / freezer systems, +27 to -30 °C

Low Temperature Refrigeration

Temperature range from about -30 to -90 °C associated with flash and biomedical freezers.

Ultra-Low Temperature Refrigeration

Generally associated with specialty systems, for example long term biological storage freezers, with temperatures in the range of -90 to -150 °C.

Cryogenic Temperature Refrigeration

Conventionally defined as any refrigeration system operating at temperatures below -160°C. Figure 1d depicts this temperature break down while noting a number of significant temperature points.

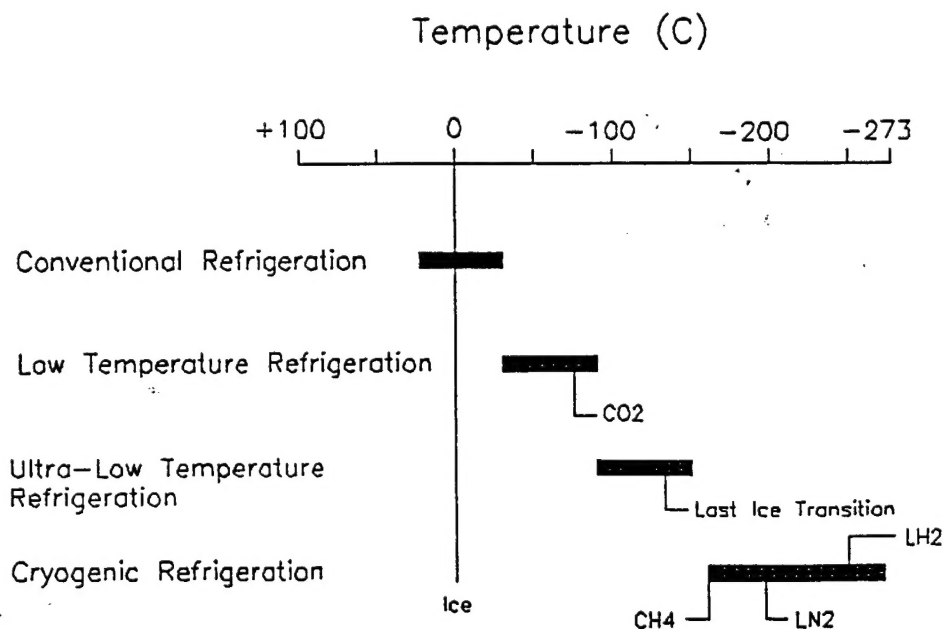


Figure 1h. Refrigeration Temperature Regimes

2.2 Candidate operational temperatures

In theory, a computer system can be operated at any temperature between ambient (or higher)

and the temperature where carrier freeze-out occurs. Realistically, this temperature range is limited by available cooling technologies, and is further limited by the efficiency of operation of cooling systems and the parasitic heat load over the temperature range. As evidenced by figure 1b above, the computing performance increases with decreased temperature, however the efficiency of cooling systems decreases with temperature as shown in figure 11h. These competing factors must be balanced to the point where the cost-benefit curves are most favorable.

The maximum efficiency of a heat engine (refrigerator) is given by the Carnot efficiency. Carnot's theorem states:

All reversible Carnot engines operating between T_{hot} and T_{cold} have the same efficiency given by $\eta = 1 - T_{cold}/T_{hot}$, and no engine operating between these two temperatures can have a higher efficiency.

In practice, a reversible Carnot engine doesn't exist, but the efficiency can be approached. In the refrigeration industry, a figure of merit of "percent of Carnot" is often used. As indicated in figure 11h, the Stirling cycle has the most consistent performance vs. the Carnot limit, with single-stage vapor compression and two-stage vapor compression providing good performance within their limited temperature ranges. When comparing the cost of implementing a refrigeration system as estimated in table II-2, we see that single and multi-stage Vapor Compression systems provide the most cost-effective performance, with Gifford-McMahon and Stirling machines a distant third and fourth. Comparing the Gifford-McMahon and Stirling performance from figure 11h leads to the selection of the Stirling cycle. On a cost basis, single stage VC machines are desirable down to -40°C , the two-stage VC coolers are useful down to -100°C and the Stirling coolers are selected for the remainder of the range, say down to -196°C (77K). Therefore, we have selected three temperature points for our cost and performance analysis:

System Terminology	Operational Temperature	Cooler Cycle Selected	COP	Speedup based on figure 1b
Ambient	25°C (298 K)	Forced-Air	N/A	0%
Refrigerated	-40°C (233 K)	Single-stage VC	1.0	39%
Low-Temperature Refrigerated	-100°C (173 K)	Two-stage VC	0.3	79%
Cryogenic	-196°C (77 K)	Stirling	0.09	245%

Table I-1 Selection of candidate Operational Temperatures

2.3 System Packaging

2.3.1 Cryogenic Package

The cryogenic package is most efficient when implemented with a vacuum enclosure or dewar to isolate the chip from the atmosphere. Atmospheric isolation is desirable to reduce or

eliminate convection heat transfer and to prevent the condensation of the various gas species on the cooled cold surface and die package. A vacuum level below 1×10^{-3} torr is required to effectively eliminate convection heat transfer. At higher pressure levels, the thermal conduction of the gas predominates until approximately 1×10^1 torr where mass-transfer effects (convection) govern.

It is not a trivial task to provide a long-life vacuum package, for even with high initial vacuum levels, the pressure will rise in a vacuum enclosure due to the outgassing of the various materials within the dewar, and the outgassing of the dewar walls themselves. Add to this the real leak rates of the joints, the diffusion rates through the dewar walls and any virtual leaks from entrapped gasses within the dewar; and the pressure levels can quickly rise to unacceptable levels. In the absence of an acceptable pressure, the cold head may not achieve the desired operating temperature.

With a large capacity cryocooler (many watts), cryopumping, or condensation of the residual gasses within the dewar may be a feasible method of "maintaining" a vacuum conducive to efficient low temperature operation, though this is not an elegant or an optimum solution. Cryopumping only works if the cold head temperature is lower than the boiling point of the liquid phase of the gas species to be pumped. The limitation of this approach is best illustrated by the experience of many dewar designers who have discovered that their cryopumped dewars cannot be shut-down because of insufficient cooler capacity to re-condense the accumulated gasses. This is also not desirable due to the problem of "wet", or condensable gasses which might cause reliability problems if allowed to precipitate on the processor die or the interconnections.

A better approach to the problem, and one developed by ICC for our IRFPA packaging technology is to select materials, design configurations and processes to minimize outgassing, diffusion and virtual leaks. This approach allows for the minimal cooler capacity necessary to overcome the parasitic heat loading due to radiation and conduction through solids and the active load of the device to be cooled. In this equation, conduction/convection through gasses is effectively neglected. Initial vacuum levels on the order of 10^{-8} torr and He leak rates on the order of 10^{-13} STD-ATM - cc/sec are typically required for a long-life cryogenic package. While the dewar design rules minimize the material and configuration options, experience shows that a myriad of possibilities exist to meet any practical need.

However, given the active heat load values shown in figure 1g above, the parasitic heat load from conduction/convection through the internal dewar medium (gas) adds a relatively small heat load. For instance, a rough-order-of-magnitude calculation yields a "convection" heat load of only 2-5 Watts. This is only 2-5% inefficiency for a 100-Watt system. There are, however additional factors to consider. If a continuous cooling load of 2-5 Watts is added to a radiation heat load of similar magnitude, a net 4-10 Watt cooling load on the package enclosure could result. Such a load will quickly lead to condensation and even frosting on the package exterior. In order to maintain thermal equilibrium, heat would have to be added to the system. Further, the "convection" heat load must be offset by additional cooling capacity. Particularly at low

temperatures, the cooler coefficient of performance (COP), or system power input needed to achieve a unit of cooling effect decreases with temperature.

To illustrate the net effect of the input power penalty due to convection cooling, we start with an assumed load of 3 Watts at -40°C and linearly increase this in proportion to the temperature difference between the ambient conditions and the device. To this we add the electrical heating which is needed to overcome the cooling effect and eliminate condensation. The result is significant, particularly at cryogenic temperatures. The graph is shown in Figure li.




Figure li – Excess wall power consumption due to a lack of a vacuum package, showing the effect of the diminishing COP at lower temperatures

In addition to the input power penalty, there is an added cost for heaters on the package exterior. We therefore conclude that a vacuum package is not strictly necessary, but offers system efficiency and potential cost advantages, particularly at the lower temperatures. A reasonable conclusion is that for device temperatures down to -100°C , a vacuum package probably doesn't make sense. Below -100°C and down to -200°C or lower, a vacuum package is warranted. We will use this two-tiered package assumption in our cost models.

Table I-x shows a comparison of using an evacuated (vacuum) enclosure vs. an insulated container.

Innovative Techniques for Cooling Computer Chips to Cryogenic Temperatures
Infrared Components Corporation
January 13, 1999 - DABT63-97-C-00721

Attribute	Insulated Technology	Complexity	Vacuum Technology	Complexity
Enclosure	Hermetic sealing using epoxies, solders or O-ring seals.	Moderate	Hermetic to 10^{-13} SCCS using ceramic-to-metal seals, solders and welds	Moderate with experience
Enclosure processing	Nitrogen purge & fill	Low	Ultra-High Vacuum (UHV) processing to 10^{-8} torr, several days	High
Internal environment stabilization	None, or simple desiccant	Low	Getter	Moderate
Insulation	Foam or double-wall "dewar".	Moderate	None - vacuum	Low, but requires above three steps
Cooling	Various, see below	Moderate	Various, see below	Moderate to high.
Chip stress management	Flexible, conductive adhesives	Low	Judicious material selection for CTE matching (materials must be vacuum compatible based on outgassing rates)	Low to Moderate with experience
Electrical Interconnects - parasitic heat load minimization	Pin connections with small area contact or individual wire bonds on a single low thermally conductive substrate	Moderate	Individual platform-to-platform wirebonds are preferred. Pins may be used, but requires a custom design	Moderate to High
Electrical Interconnects - vacuum sealing	Epoxy potting of pins or low-hermeticity ceramic to metal seals	Moderate	Ceramic-to-metal seals	Moderate to High
Heat Transport System	Conduction through copper or other high thermally conductive material and/or working fluid phase change	Low to High	Conduction through ceramics, silicon, copper or other high-thermally conductive material and/or via a phase-change heat pipe or possibly through immersion in a cooled-fluid bath	Low to High
Overall Complexity		Moderate		Moderate to High

Table I-2 Comparison of relative complexity for vacuum vs. insulated enclosures.

2.2.2 Refrigeration Package

The packaging considerations for the nominal -40°C system are much less rigorous from a design and processing perspective. The chief objectives in the enclosure design is the elimination of wet gasses such as water vapor which might condense on the chip and cause damage or undesirable operation. A second objective is related to the first, which is to keep the case temperature below the dew point. While this can be achieved by artificially maintaining the dew point at very low levels (this might make sense for larger parallel processing

computers), a more practical approach is to thermally isolate the package exterior from the cold internals. A further benefit to thermal isolation is the reduction in parasitic heat load on the cooler, though this factor is less efficiency leveraging in the Low Temperature Refrigeration system than it is in the cryogenic system due to the higher COP in the higher temperature ranges. A third option is the addition of heaters at the package exterior, which serves to maintain the temperature above the dew point. This last option requires additional parts and will result in additional heat load on the cooler (due to higher external ambient) unless the case temperature is monitored and the heater power is regulated using a closed-loop system. KryoTech uses this approach on their cooled computer systems. Heaters may be necessary at the interconnect location due to the high (tens of watts) of parasitic heat load at this point.

2.2.3 Ultra-Low Temperature Refrigeration

Systems with operating temperatures below -90°C down to approximately -150°C have many of the same characteristics as the Cryogenic system, and will benefit from the use of an evacuated enclosure. Since the COP logarithmically decreases with decreased temperature (see Figure 11h), more system efficiency is gained with the lower temperatures. In the case of a single or multi-stage vapor compression system operating at -100°C , the COP is high enough that the potential efficiency gains of a high-vacuum enclosure (under 10^{-3} torr) may be offset by the higher cost and complexity of such an enclosure. See figure 1i. Still, a modestly hermetic package (perhaps to 10^{-6} STD-ATM - cc/sec) which seals out the "wet" gasses and which insulates the cold components from the warm ambient conditions is required for the "non-vacuum" package.

2.3 Package Thermal Budget

The thermal management of a sub-ambient cooled system is different than a traditional higher-than-ambient system. This is due to the orientation of the thermal gradients for these two systems. For example, the thermal gradient of a higher-than-ambient system is positive whereas a sub-ambient system has a negative thermal gradient. Figures 1e and 1f show the thermal balance schematically.

The heat balance equations for the control volume of the chip package for a higher-than-ambient system is:

$$q_{\text{generated}} = q_{\text{cooling}} + q_{\text{convection}} + q_{\text{radiation}} + q_{\text{conduction}}$$

As is shown schematically in Figure 1j.




Figure 1j - Schematic view of heat transfer modes, higher than ambient system

A sub-ambient system has reversed direction of heat flow for the convection, radiation and conduction terms:

$$q_{\text{generated}} = q_{\text{cooling}} + (-q_{\text{convection}}) + (-q_{\text{radiation}}) + (-q_{\text{conduction}})$$

Rearranging the latter expression for a sub-ambient system, we have:

$$q_{\text{cooling}} = q_{\text{generated}} + q_{\text{convection}} + q_{\text{radiation}} + q_{\text{conduction}}$$

A schematic depiction of the heat transfer modes is shown in Figure 1j.




Figure 1k - Schematic view of heat transfer modes, sub-ambient system

In a higher-than-ambient system, the convection, radiation and conduction terms are beneficial as they supplement the cooling system. In a sub-ambient system, these terms provide an additional load on the cooling system, and hence are parasitic in nature. This fact requires a paradigm shift in electronics packaging.

Quantifying the relative values of the parasitic heat loads, we perform some first-order calculations that yield the following:

2.3.1 Radiation

A preliminary calculation, places the value of radiation heat transfer in the chip package at 3 watts. Providing a design margin, we will use 5 watts as the radiation heat load for thermal budgeting purposes. The actual radiation heat load in the final configuration will be higher or lower depending on geometric configurations, surface areas and surface conditions. This estimate is valid at or below -100°C .

2.3.2 Convection heat transfer

Convection heat transfer has been estimated as 3 Watts at -40°C , but linearly decreases with temperature to a value of approximately 10 Watts at 77 Kelvins (-196°C).

2.3.3 Conduction along interconnections

As calculated in section 3.0, the parasitic heat load on typical phosphor-bronze pins (.018" diameter) is approximately 0.450 Watts per pin for a 180°C thermal gradient. This gives rise to a parasitic heat load of 185 Watts for 413 pins.

Preliminary calculations show that an improvement of 5-10X in the thermal loss through the pins can be achieved while maintaining good electrical performance. Such a system would entail the use of an intermediary section of conductor to provide the thermal resistance necessary. Using a five fold reduction as a design baseline, the per pin parasitic heat load becomes 0.09 Watts. Based on 413 pins, this would be 37 Watts.

2.3.4 Total Thermal Budget

In the absence of achievements in the area of reduction of thermal loads due to the electrical interconnections, the parasitic load on the pins for a -100°C system becomes a large percentage ($>80\%$) of the overall cooling load, requiring large cooling capacities to overcome. This is shown graphically in Figure 1m.




Figure 1m - Cooling Load at -100°C , standard pins




Figure 1n - Cooling Load at -100°C for 5x improvement in pin parasitic thermal load




Figure 1p - Cooling Load at -100°C for 10x improvement in pin parasitic thermal load

If we assume that the pin parasitic load can be reduced by a factor of 10, we nominally size a -100°C cooler at 100 Watts of net cooling capacity for a current application such as a DEC Alpha processor. A cooler with future legacy should have at least 150 Watts capacity, and perhaps 200 Watts for a five-year product life.

2.4 Thermal Load Summary

The following assumptions have been made for calculation of the thermal load summary:

1. The chip heat dissipation is constant over the temperature range, and is estimated to be 60W.
2. The number of interconnections is 500, and the I/O physical configuration is an Interstitial Pin-Grid Array (IPGA). A 5x reduction in pin heat transfer has been applied.
3. High-performance insulation and inert-gas fill are used in the package interior for all systems as a worst-case analysis.
4. Exterior heaters are used to eliminate condensation on the package exterior.

System	Radiation	Pin load *no reduction taken	Convection	Total thermal load (with 60W active)	Total electrical load from cooling	Total power consumption with package exterior heating and system cooling
Ambient	N/A	N/A	N/A	60W	5W	5W
Refrigeration	1.6W	86.5W*	3W	151.1W	75W	166.1W
Low-Temp Refrigeration	2.3W	32.8W	6W	101.1W	337W	378.1W
Cryogenic	2.6W	57.7W	10W	130.3W	1447W	1517.3W

Table I-3 – Thermal load summary matrix

As shown in table I-3, the parasitic heat load increases inversely to operational temperature.

3.0 Electrical interconnection

The electrical interconnections between the processor and the motherboard (or bus) are typically accomplished using any of three technologies:

- Pin Grid Array (PGA) or Interstitial Pin Grid Array (IPGA), such as the DEC Alpha processor.
- Card edge connection to the bus, such as Intel Pentium®II "Socket-1".
- Ball Grid Array (BGA), such as the IBM PowerPC processors.

Of the signals on these connection systems, there are three different types of electrical data that comes out of these systems:

- Signal leads comprising high-speed digital transmission lines for which the reduction of capacitance and inductance is paramount.
- Power leads comprising low-resistance, low inductance DC conductors @ about 3.3VDC @ >1A
- Ground leads comprising low-resistance DC conductors @ about 3.3VDC @ >1A

For a typical processor, the DEC Alpha 21264, the pin count and impedance considerations are as follows:

Number of Pins	Signal Type	Impedance Consideration
264	High-speed Digital	Minimize Capacitance, Inductance
74	Ground	Minimize Resistance
46	External Power (Vdd)	Minimize Resistance, Inductance
22	Internal Power (Vddi)	Minimize Resistance, Inductance
7	Unused	possibly eliminate to reduce parasitic heat load

The performance of a cryogenic electrical connection is measured in both electrical resistance (or impedance) and thermal conductance. The fundamental objective of the connection is to provide a high thermal resistance (low conductance) and low electrical resistance (impedance) path. These goals are competing, as is best evidenced by the generalized ratio of these two properties, which is known as the Wiedemann-Franz law:

$$k/sT = \text{Lorentz Constant} \approx 2.45 \times 10^{-8} \text{ W} \cdot \text{W/K}^2$$

Where k is the thermal conductivity and s is the electrical conductivity. T is the temperature in Kelvins.

For the case of electrical resistance, the Wiedemann-Franz law implies that no thermal advantage can be gained by reducing electrical resistance as the two properties are related.

Innovative Techniques for Cooling Computer Chips to Cryogenic Temperatures
Infrared Components Corporation
January 13, 1999 - DABT63-97-C-00721

This relationship holds true for metals at temperatures of several hundred Kelvins, below this range the "law" breaks down. In the extreme case of a superconductor, the proportionality "constant" is 0. In the first-order analysis, we will assume that the Wiedemann-Franz law applies, in the second-order analysis, we will seek materials, particularly alloys that will provide advantage.

Mechanical strength (such as tensile strength) and perhaps fatigue resistance must also be considered in selecting a material. A second order analysis should also consider potential Thomson, Seebeck or Peltier Effects.

In the case of signal leads, the electrical resistance must be modeled as Impedance, which is the sum of resistance, capacitance and inductance terms. In the first-order analysis, we will use a "lumped element" model which is valid for frequencies up to 100 MHz for short conductor lengths.[1] The second-order analysis will incorporate transmission line theory.

The thermal load on the cooling system due to the pins of a PGA package (such as the DEC Alpha) is as follows:

$L = .00254$ $\Delta T = 330\text{-K} - 150$	Effective (active) length of pins in meters Thermal Gradient dropped along the pin length.
---	---

Thermal conductivity of Phosphor-Bronze Pin, average over the temperature range.

$\text{diameter} = .001\text{-in.}, .002\text{-in.}, .018\text{-in.}$ $r(\text{diameter}) = \frac{\text{diameter} \cdot 25.4 \cdot 10^{-3}}{2 \text{ in}}$ $A(\text{diameter}) = \pi \cdot r(\text{diameter})^2$ $q_{\text{pins}}(\text{diameter}) = k \cdot A(\text{diameter}) \cdot \frac{\Delta T}{L}$	Diameter of pin in inches Radius of pin in meters, calculated from diameter. Cross-sectional area of pin in meters ² Conduction heat transfer for one pin
--	---




Figure 1q – parasitic heat load on a cooled processor, thermal gradient between 330K and 150K

For the DEC Alpha 21164PC processor, there are 413 pins, and assuming an active length of 0.1 inches, the total heat load is:

$$\text{diameter} = .018 \text{ in}$$

$$q_{\text{pins_total}} = q_{\text{pins}}(\text{diameter}) \cdot 413$$

$$q_{\text{pins_total}} = 192.199 \text{ kg} \cdot \text{m}^2 \cdot \text{sec}^{-3} \quad [\text{Watts}]$$

The capacitance of two parallel conductors is a function of the diameter, length and distance between the (termed "pitch" for a conductor array). While the capacitance of an array is more complex, this two-conductor approximation is useful for determining the scope of the problem. We therefore can plot capacitance as a function of heat load in Watts for a 50-mil interstitial pin pitch.

$$\text{pitch}_{\text{pins}} = .05 \text{ in} = 2.54 \cdot 10^{-2} \frac{\text{m}}{\text{in}}$$

$$q = .0001 \text{ watt}, .002 \text{ watt} \dots .5 \text{ watt}$$

$$\text{Capacitance}(q) = \frac{\pi \cdot 9 \cdot 10^{12} \cdot L}{\text{acosh} \left[\text{pitch}_{\text{pins}} \cdot 0.5 \cdot \sqrt{\frac{k \cdot \pi \cdot \Delta T}{q \cdot L}} \right]}$$

Innovative Techniques for Cooling Computer Chips to Cryogenic Temperatures
 Infrared Components Corporation
 January 13, 1999 - DABT63-97-C-00721




Figure 1r – Lead to Lead capacitance decreases with decreased heat load

The Lead/Lead Inductance can be similarly calculated vs. heat flow:

$$\mu = 4 \cdot \pi \cdot 10^{-7} \text{ H/m}, \text{ Inductance (q)} = \frac{\mu}{\pi} \cdot \text{acosh} \left(\frac{0.5 \cdot \text{pitch}}{\text{pins}} \right) \cdot \frac{k \cdot \pi \cdot \Delta T}{q \cdot L}$$

This is displayed graphically in figure 1s.




Figure 1s – Lead-to-Lead Inductance increases as pin heat load is minimized.

Innovative Techniques for Cooling Computer Chips to Cryogenic Temperatures
Infrared Components Corporation
January 13, 1999 - DABT63-97-C-00721




Figure 1t – Schematic of interconnect thermal socket

The above suggests that thermally optimized electrical interconnections can be achieved by the simple expediency of thinning the pins within a matrix of pins as shown in figures 1t and 1u.




Figure 1u – Detail of area of pin diameter reduction

A pin matrix such as is shown above provides the strength necessary for reliable fabrication and use. A review of the currently available pin diameters indicates minimum diameters of 0.012 inch for commercial off-the-shelf sockets and pins. However, preliminary calculations indicate that pins as small as 0.005 inch can be implemented in a matrix configuration. Calculations for strength indicate that for a stainless steel pin with 0.1 inches of "free" length, compressive stress is the limiting factor for axial load resistance above 0.005 inches, and Euler buckling limits the load below that point. The calculations indicate that a matrix of 413 5-mil diameter stainless steel pins has an axial load resistance of nearly 150 pounds, which allows for an ample factor of safety axially. Load resistance orthogonal to the pins is however unacceptable at only 1.5 pounds for same matrix. This indicates that additional support is needed to resist shear forces on the assembly, which will of course add to the

thermal load. Another strategy is to bias the pin geometry towards greater geometric rigidity (higher moment of inertia of the cross section). This can be accomplished through the use of hollow pins to permit a more efficient use of both electrical conductance area and mechanical strength. Alternatively, non-prismatic pin geometry that is thinner in the center of the span than on either end permits efficient use of cross sectional area in those areas with the highest bending moments.




Figure Iv – Axial load resistance of a stainless steel pin, 0.1 inches in length




Figure Iv – Maximum transverse load on a stainless steel pin, 0.1 inches long

4.0 Chip stress management

Thermally induced stresses are present in the die-substrate system due to the difference between the stress-free temperature and the "final" temperature whether transient, static in-use or during storage. This phenomenon takes three principle forms, all related to the following well-known equation:

$$\Delta L = \alpha L_0 \Delta T$$

Where:

- ΔL = change in length
- L_0 = initial length
- α = Coefficient of Thermal Expansion (also CTE)
- ΔT = change in temperature

This equation can be used to provide input into any of a number of well known expressions relating stress to deformation.

The most obvious form is the "bi-metallic" stress in a static state, which is when the system has achieved equilibrium at some temperature different from the stress-free point. In this case, the stresses can be calculated from knowledge of α and the stress free and final temperatures.

The second form of thermally induced stress is strongly related to the first except that a knowledge of the temperature variation in the CTE can be employed to discover critical temperature points where the deviation of the CTEs of the different materials reach a local maximum.

Finally, thermally induced stresses can take a dynamic form caused by uneven cooling. This is due to the change length of one finite element as compared to an adjacent element, and is caused by CTE and/or temperature gradients. This occurs in homogeneous materials that are unidirectionally or otherwise unevenly cooled, and is greater when geometric and/or material property differences occur along the direction of heat transfer. This phenomenon is difficult to predict and requires transient heat transfer and simultaneous structural analysis. Finite Element Analysis (FEA) is typically employed for this purpose, but it is difficult to model and even harder to verify experimentally. Experimental verification is typically performed optically, using techniques such as Moire Interferometry.

The principle differences in the problem of chip stress management in the Low Temperature Refrigeration and Cryogenic systems are the magnitude of the problem (cryogenic is more problematic) and the breadth of available means to manage the effects. For a cryogenic system which has a room-temperature stress free temperature and a 77K operating point, the ΔT is 220°C. For a system at -40°C, the ΔT is 70°C, or only one third of the cryogenic value. The Refrigerated system designer enjoys considerably more freedom in developing stress management techniques due to the greater variety of acceptable materials allowed

Innovative Techniques for Cooling Computer Chips to Cryogenic Temperatures
Infrared Components Corporation
January 13, 1999 - DABT63-97-C-00721

by the relaxed out-gassing and low-temperature endurance requirements (most traditional materials can be used). The cryogenic system designer must contend with out-gassing issues, higher temperature processing (may be higher than the stress-free value) necessary to de-gas the system, and a limited pallet of acceptable materials.

A number of very acceptable adhesives and substrate materials have been developed for cryogenic packaging for the IR industry, and most of the typical palette of ceramic and metals may be used to good effect. Experience shows that silicon devices of typical processor size may be packaged at cryogenic temperatures with little risk and need for complexity. The principle increased complexity that is warranted for cryogenic stress management is at the analysis phase, and not necessarily in production.

ICC has considerable experience with the analysis and design of stress management techniques, and therefore has not performed any analysis of thermally induced stresses under this contract.

Section II - Refrigeration Technology Review and Cycle Selection

Introduction

In selecting refrigeration technology for use in the cooled computer system, we first cast a wide "net" over all possible choices and then down-select to two systems that show most promise, and provide a more detailed comparison of these two systems.

1.0 Refrigeration Systems

A wide variety of thermodynamic cycles can provide the desired refrigeration effect over the temperature range of interest in cooled computer systems. The fundamental cycles are shown in Figure 11a with a number of specific cycles noted. Each cycle generally has an operating temperature range where its cooling performance is maximized. At the same time there are usually temperature limits associated with each cycle where the cycle ceases to operate or its performance is so poor that it will not provide an acceptable cooling capacity. While a number of "rating" factors are used to define the relative performance of a refrigeration system, the cycle Coefficient of Performance (COP) will be used throughout the current effort. The COP is defined as the ratio of cooling effect to the power required. Since both the cooling capacity and the input electric input power are measured in watts, COP is a dimensionless parameter. It is assumed that the cycles of interest will be electrically driven by the same basic energy source which is providing power to the computer system. Since the COP of all refrigeration systems fall dramatically as the operating temperature is lowered, the actual power required to produce a fixed amount of cooling rapidly increases. The ideal refrigeration cycle COP is defined as that of a Carnot device operating between the lowest (T_{cold}) and highest (T_{hot}) temperatures present. The COP of this ideal cycle is simply $T_{cold}/(T_{hot}-T_{cold})$. In a number of cases it is useful to rate a cycle as the ratio the COP of the actual cycle to that of a Carnot cycle COP thus defining a "% of Carnot" for the particular cycle. This approach is quite commonly employed in the cryogenic temperature regime.

For the chemical refrigeration systems noted below, the restriction to electric input power will require conversion of electrical energy into thermal energy (via an electric heater) since these cycles are "driven" by thermal rather than mechanical energy input. An additional restriction placed on the refrigeration cycles under consideration is that they be "closed" so that the same working fluid is employed throughout the cycle process without being replenished from outside sources. These latter "open" systems are characterized by either the use of air as a working fluid (limited temperature range) or the boil off of a low temperature fluid such as liquid Nitrogen.




Figure IIa - Refrigeration System Cycles

A refrigeration system also exists which is based on the unique properties of certain solid state materials that provide a cooling effect when electrical energy is supplied. The most common version of this type of refrigerator is the Thermoelectric cooler. While these types of coolers possess a number of operational advantages such as accurate temperature control and no moving parts, their COP's are typically well below the thermodynamic cycles. This low COP implies high electrical input power levels (with associated high heat rejection rates) that will effectively limit these coolers to relatively low cooling capacity applications. For this reason the Thermoelectric cooler will not be considered in the current program.

1.1 Closed Cycle, Phase Change, Vapor Compression Cycles

In the vast majority of applications involving the conventional and low temperature refrigeration regimes, the vapor compression or reversed Rankine cycle is used to provide the necessary refrigeration capability. This type of refrigeration system consists of a "cold" evaporator, gas compressor, "warm" condenser, and a refrigerant control valve or expander. The refrigerant working fluid must change from a gaseous to liquid state over the temperature range defined by the condenser and evaporator and as such plays a dominant role in determining both cycles thermal performance and the operating temperature limits. While it is not possible to set absolute minimum temperatures limits for the various vapor compression system refrigerants, it is possible to use compressor inlet pressure constraints as a way to define some practical lower limits for these systems. Table II-1 lists a number of conventional refrigerants along with the temperature equating to a gas pressure of .1 MPa (14.5 psi) a value likely to be below that employed in practical compressor inlet valve operation. Not listed are a number of proprietary refrigerant mixtures which are used to extend the operating temperature range of single and "cascaded" conventional vapor compression cycles. It is important to note that a number of the refrigerants noted have been banned due to ozone depletion characteristics (for example R-12). In some cases effective substitutes have been found (for example R-134a for R-12) but in other cases refrigerant mixtures containing combustible mixtures are employed as the replacement. A number of safety issues will

likely have to be addressed before these fluids are used in the US marketplace.

Table II-1 Vapor Compression System Refrigerants

Refrigerant	Temperature Limit, °C
R-11	23.7
R-12	-30
R-22	-41
R-134a	-26
R-500	-33
R-502	-45
R-503	-89
R-504	-57
R-1150	-104
R-14	-128

The lower temperature limits of various modified cycles, described in a later part of this section, cannot be defined in this manner since the mixed refrigerant cycles depend on a number of heat exchanger / gas separation stages prior to the final cold expansion stage. The actual remaining gas which is expanded in the cold head are usually mixtures of nitrogen and other gases such as argon, neon, or a number of combustible hydrocarbons.

For lower temperature operation it is common to "stage" or "cascade" the simple vapor compression cycle by adding additional stages of compression and various evaporator / condenser combinations. With carefully selected refrigerant mixtures, the 2-stage arrangement is used routinely down to approximately -90°C to support the biomedical industry. Incorporation of another stage yields a 3-stage system which, when combined with unique mixed refrigerant combinations, make it is possible to attain temperature down to -140 to -150°C range. While these systems retain all the advantages and disadvantages of the classic single stage system described above, they are also plagued by problems of the oil from the compressor migrating to the colder portions of the system and clogging the critical expansion valve.

In an attempt to overcome the problems encountered with the low temperature staged vapor compression cycles, a number of cycle variants have been developed. The majority of these "cycles" incorporate unique fluid mixtures, an expander valve in the last "cooling" stage, and a number of heat exchangers which are used to separate the various components of the mixture as they liquefy. MMR's "Klimenko" cycle and the

ADP's "CryoTiger" are representative of these devices. Both are characterized by lower temperature limits, on the order of -195°C , and the use of refrigerant mixtures which contain flammable components. While more efficient than conventional staged systems, even these advanced cycles have relatively poor thermal performance and can degrade due to the refrigerant mixture "separating" at the wrong temperatures or carry over of compressor lubrication into the colder portions of the system leading to excessive contamination.

Due to its mechanical simplicity and good thermal performance, the basic vapor compression refrigeration system has been built in phenomenal numbers and represents an extremely low cost approach to system cooling. Additional advantages and disadvantages of the basic cycle are noted below:

Advantages

1. Compressor can be located at significant distances from the cold evaporator with only minor impact on the overall system performance.
2. The refrigerant present in the evaporator provides extremely high heat transfer coefficients ($> 20,000 \text{ W/m}^2\text{C}$) which simplify device cooling.
3. Simple scaling to a wide range of cooling capacities.
4. Compact mechanical configuration.
5. Low cost per watt of cooling down to approximately -90°C , rapidly increasing costs at lower temperatures.

Disadvantages

1. Operating temperature range entirely defined by the refrigerant fluid's thermophysical properties.
2. Limited number of acceptable refrigerants (non-toxic, non-flammable, environmentally approved) are available, particularly at the lower temperatures of interest.
3. Become mechanically complex when configured for low temperature operation.
4. COP falls dramatically at lower temperatures - high electrical input power required.

1.2 Closed Cycle, Single Phase (Gas) Cycles

The gas cycle family of refrigeration systems by definition employ a gaseous working fluid throughout the cycle. A minor exception is the Joule Thomson cycle in which throttling of the gas across the expander valve can lead to the formation of a gas / vapor mixture. In the case of a JT refrigerator, any liquid within the vapor is evaporated by the device being cooled as rapidly as it is being formed and all the refrigerant returns in the gas phase to the compressor.

1.2.1 Regenerative

Regenerative cycles employ a heat exchanger / thermal storage device (regenerator) which is used to alternately cool and heat the gas as it flows between the hot and cold

spaces within the cycle. In this family of cycles the same gas flows through the same regenerative matrix in each direction during each "hot" and "cold" flow period which occur once during the cycle. While these cycles will operate without a regenerator, the thermal performance is so low that they provide little or no net cooling effect at the temperatures of interest. A significant advantage of all of the gas cycles is the fact that the majority of the gaseous refrigerants do not have the thermophysical constraints such as those which severely limit the operating temperature range of refrigerants for the vapor compression cycle.

1.2.1.1 Stirling Cycle

The Stirling cycle represents one of the oldest gas cooling cycles. While initially developed as a heat engine (thermal energy in - mechanical power out) in the early 1800's, by the mid 1800's it had been discovered that it would operate as a refrigeration system by driving the system via an external mechanical power supply. The basic cycle employs two moving parts, a power piston and a displacer which can be driven via a mechanical mechanism (kinematic drive) or utilize the Free-Piston concept in which the components operate at a near resonant condition without the need for a mechanical drive mechanism. The expander portion of the system contains the displacer and cold head where the actual cooling effect is produced. The gas compressor portion of the system contains the power piston, the drive motor, and a heat exchanger which dumps the cycle waste heat to the surroundings. Current Stirling cycle coolers have operated over the entire temperature range of interest ("air conditioning" to deep cryogenic temperatures) with cooling capacities up to the KW level in some industrial sized systems. The advent of the Free-Piston concept has allowed the use of non-lubricated drive systems which eliminate the contamination problems associated with the lubricant. The use of non-contacting bearing systems (gas or flexure based) have radically reduced the potential of wear induced failures.

Advantages

1. Well established technology base.
2. Basic cycle provides good cooling performance over the entire temperature range of interest.
3. Very intense cooling effect occurs at the cold head.
4. Can be easily scaled to high capacities (for example units with capacities up to 4,000 watts at -200 C are commercially available).

Disadvantages

1. For best cooling performance the compressor and cold head should be as close together as possible.
2. The cold head can be quite small relative to the object being cooled so that some form of conductive heat "spreader" or heat transport system may be required.
3. Relatively high cost per watt of cooling.

1.2.1.2 Pulse Tube

The pulse tube can be described loosely as a variant of the Stirling cycle in which the

moving displacer has been removed and replaced with an empty tube, a gas orifice, and a gas reservoir. These latter components provide the proper phasing between the cyclic mass flows and pressure cycling to provide a net refrigeration effect on the cold end of the device. Thermodynamically the Pulse Tube will always have lower efficiency than a comparable Stirling cycle refrigerator. Current Pulse Tube developments have focused on low capacity units operating at temperatures at or below 80 K. Under these conditions the fundamental thermodynamic advantages of the Stirling are quite small with the simplicity of the Pulse Tube in many cases will make it the refrigeration cycle of choice. At higher cooling capacities (>10 watts) and higher temperatures (particularly above -150°C) the relative performance of the Pulse Tube falls dramatically compared to the Stirling refrigerator. This effect is shown in Figure 11b, note that this is a comparison of ideal Stirling and Pulse tube cycles so a number of losses that will narrow or eliminate the performance differences at cryogenic temperatures are not indicated. The lower performance of the Pulse Tube at higher temperatures however is clearly evident.

The basic Pulse Tube concept can also be employed with valved compressor systems, see GM refrigerator section.




Figure C. Relative Comparison of Pulse Tube and Stirling Cooler Performance vs Cold Head Temperature

Advantages

1. Elimination of the moving displacer - improved reliability.
2. Intense cooling effect in cold head.

Disadvantages

1. Lower performance, particularly at higher capacities and operating temperatures.
2. Compressor and Pulse Tube must be located close together to minimize performance losses.
3. Higher performing Pulse Tubes are generally of the "in-line" configuration which places the cold head in the center of the device rather than on the end - packaging of the device can be a problem.
4. Higher performing Pulse Tubes tend to operate at lower (<30 Hz) frequencies which can result in physically large drive motors.
5. Scaling to large sizes (>100 watts capacity) needs confirmation.
6. Fundamentally about the same cost as Stirling coolers - higher cost per watt of cooling.

1.2.1.3 Vullemier (VM) Refrigerator

The VM refrigerator is a heat activated systems in which a thermally driven gas compressor drives a gas cycle refrigeration system. Both the hot "drive" side of the system and the "cold" refrigeration side of the system share a common gas supply. Since the cycle is driven by the temperature-difference-induced pressure changes, there is no need for an electrical drive motor to power the cycle. VM systems found considerable use in the early development of small cryocoolers for spacecraft applications since the electric drive motor and mechanical mechanism of conventional Stirling cycle coolers were not required. The VM cycle is characterized by good thermal performance at low temperatures but, because it usually operates at low frequencies (<30 Hz), is physically large.

Advantages

1. Good cooling performance at low temperatures.
2. Intense cooling at cold head.
3. Low vibration.
4. Heat activated (by electric heaters) eliminating potential motor EMI.
5. Easily scaled to large sizes.

Disadvantages

1. Entire system must be located close to cooled device.
2. Physically large.
3. For good performance, high heater temperatures are required (>650°C) which can lead to operating life problems.

4. No established end users - very high cost per watt of cooling.

1.2.1.4 Gifford-MacMahon (GM) and Derivative Cycles

The GM cycle is characterized by the use of valves to properly sequence the pressure cycling of the refrigerator. Due to the use of valves, the actual cycle pressure ratio (P_{\max} / P_{\min}) is significantly greater than that of other gas cycle refrigeration systems. This high pressure ratio compensates for the low operating speeds of the cold head, managing to keep the physical size of the cold head similar to Stirling cycle equipment of equal cooling capacity. The gas compressors employed in the GM units are conventional vapor compression system compressors with only minor mechanical modifications. A moving displacer and regenerative heat exchanger, very similar to that of a Stirling cycle refrigerator, are located in the expander / cold head assembly. The motion of the displacer is controlled by the action of the valves. The gas compressor employs lubricants which must be removed from the gas stream before it reaches the cold head. This is generally done via the use of molecular sieves / filters which do require periodic replacement. The basic cycle variants can be operated from liquid Helium temperatures (-270°C) up to approximately -75°C .

Due to its relatively low performance, a number of GM manufacturers are investigating substituting a Pulse Tube for the conventional GM cold head. Not only would this lead to a performance improvement, but also the elimination of the moving displacer assembly located in the expander assembly. It is important to note that the same performance disadvantages at higher operating temperatures also exist for the GM based Pulse Tube.

The GM refrigeration system has established an extensive user base, particularly in the electronics manufacturing area where it is utilized in high performance cryopumps that are necessary for ultra-clean vacuum systems.

Advantages

1. Very well established technology and manufacturing base.
2. Can be scaled to higher capacities, 100s of watts.
3. Cold head can be located at significant distances (10s of meters) from the compressor without significant loss in performance.
4. The gas compressor portion of the system is based on commercial vapor compression refrigeration systems - low cost.
5. Motor induced EMI can be eliminated by remote location relative to cold head.
6. Potential performance improvements via the use of a Pulse Tube.

Disadvantages

1. Significantly lower thermal performance than the Stirling cycle, systems are large and noisy.
2. Generally operate at low frequencies ($<10\text{ Hz}$).

1.2.2 Recuperative

The recuperative cycles are similar to the regenerative cycles in that a gaseous medium is used as a working fluid; however in these cycles the gas flows between the warm and cold components are separated into individual flow passages and the actual storage of thermal energy within the heat recuperator is minimized.

1.2.2.1 Brayton Cycle

The Brayton cycle simply acts as a throttling device in which high pressure gas flows through an expansion device (turbine) exiting at low temperature which can be used for cooling. The power derived from the expansion of the gas through the turbine can be used to assist in driving the required gas compressor. In most low temperature systems, the gas is cooled prior to reaching the expansion turbine by the cold gas which is flowing back to the compressor after providing the cooling effect. The recuperative heat exchanger provides this function.

Advantages

1. The use of high speed turbo machinery in Brayton systems provides very compact systems with high power to volume and weight ratios.
2. Can be easily scaled to high power levels.
3. Various non contacting bearing schemes are available for long life systems.

Disadvantages

1. The recuperative heat exchanger (cold box) effectiveness is critical to performance, can lead to large and complex recuperator.
2. Does not scale well to small sizes due to losses which are characteristic of rotating equipment.

1.2.2.2 Joule Thompson (JT) Cycle

The JT cycle represents a composite cycle having a number of characteristics of the more advanced vapor compression cycles incorporating intermediate heat exchangers. The fundamental JT refrigerator employs the JT effect which causes a gas to cool when throttled through an expander valve. The gas refrigerant is compressed, passes through an after cooler to remove the heat of compression, flows through the recuperator heat exchanger where it is cooled by the cold gas returning from the cold head. The cold high pressure gas is expanded by passing through a throttle valve (orifice). The cold gas, possibly containing liquefied gas in the form of a vapor, provides the cooling effect at the device. After providing the cooling effect this low pressure gas returns to the compressor via the recuperative heat exchanger. Most gases can be employed in this manner to reach their normal boiling point, for example -196°C for Nitrogen. For a number of gases with lower boiling points (for example Hydrogen at -250°C) the JT inversion temperatures are well below ambient conditions thus requiring that the gas be pre-cooled by some form of an auxiliary refrigeration system.

Since all of the expansion work is lost to the cycle, the JT systems have very low

thermal performance in comparison to other cycles noted above.

Advantages

1. Very simple cycle.
2. Cold head can be located far from the compressor without significant losses in performance.
3. Very intense cooling at cold head.
4. Essentially no cold head vibration.

Disadvantages

1. Low thermal performance - high input power requirements, systems become large and noisy.
2. Large and expensive recuperative heat exchangers are needed.
3. Expander (throttle) valve prone to clogging due to contamination.

1.3 Chemical

The absorption refrigeration cycle is thermally (heat) activated and requires the use of an absorbent and refrigerant working fluid. The cooling load causes an increase in the concentration of the refrigerant in the absorbent forming a refrigerant-rich mixture. Thermal energy in the form of heat is used in the boiler portion of the system to boil the refrigerant-rich mixture essentially driving off the excess refrigerant. The refrigerant-poor mixture is then recycled to the absorber (cold) portion of the system where it again absorbs the refrigerant that was driven off in the boiling process. While a number of fluids have been utilized in absorption refrigeration systems, the Lithium Bromide and Ammonia types represent the highest level of development.

The refrigerant fluid totally defines the operating limits of the absorption refrigeration cycle. Since water is the actual refrigerant in the Lithium Bromide system the lowest temperatures possible are on the order of +4°C. The Ammonia systems lowest temperature limit is on the order of -60°C.

Advantages

1. Excellent heat transport properties of their working fluids.
2. Ability to locate the boiler at significant distances from the (cold) absorber without impacting performance.
3. Minimal number (if any) of moving components required.
4. Low cost per watt of cooling.

Disadvantages

1. Limited operating temperature range.
2. Ammonia system uses a toxic and flammable working fluid.

2.0 Impact of Heat Transport System Characteristics

All of the refrigeration systems discussed above will require some form of a "heat transport system" which couples the device being cooled to the cycle's refrigerant as shown schematically in Figure IIc. In some cases the heat transport system can be a simple solid wall of a high thermal conductivity material such as copper which separates the device from the refrigerant. In other cases it may have to be a complex thermal path made up of solid materials, fluids which undergo a phase change, and the refrigerant. The selected heat transport system must take into account the end users system requirements, for example ease of cooler or device replacement as well as the fundamental characteristics of the refrigeration system utilized.




Figure IIc. Heat Transport System

The basic performance characteristic of the heat transport system is its thermal resistance (generally expressed in term of temperature difference per unit energy flow, for example $^{\circ}\text{C}/\text{watt}$). The goal is to have as low a thermal resistance as possible so as to minimize the temperature that the refrigerator must operate **BELOW** the desired temperature of the cooled device. We note that the computer system performance and cooler sizing should also take into account the thermal resistance of the device package. For instance, a typical computer processor has a 60°C thermal resistance from chip to package heat rejection surface of 0.2 to $0.4^{\circ}\text{C}/\text{W}$. However, since performance figures of processor speed improvement vs. temperature are given in terms of external case heat rejection temperatures, we will not factor this resistance into the overall system thermal budget. However, it should also be noted that the thermal resistance of a composite chip package will vary with decreased temperature due to the temperature-dependant thermal conductivities of these materials in the temperature region of interest.

A first order estimate of the impact of thermal resistance on performance of a refrigeration system is shown in Figure IId. For this particular case a gas cycle refrigerator (Stirling) operating at a fixed fraction of Carnot (assumed constant at 20%) is used for cooling a device which dissipates a thermal load of 20 watts. The system rejects the cycle waste heat to a 30°C ambient. The approximate input power required to drive the refrigeration system is shown as a function of thermal resistance for three different device temperatures (-100, -150, and -200°C).




Figure IId. Impact of Heat Transport System Thermal Resistance

As can be seen thermal resistance below approximately $.1^{\circ}\text{C/W}$ do not play a significant role in defining system COP or required input power. In the range of $.1$ to $.4^{\circ}\text{C/W}$ the impact on power requirements becomes more pronounced particularly at the lower device temperatures.

Figure IId also graphically depicts the relative impact of decreasing device temperature on input power requirements (or COP's), independent of the actual thermal resistance of the heat transport system.

3.0 Heat Transport System Component Properties

Cooling of the actual device or CPU will require the removal of a modest amount of heat

from an extremely small area. The resulting surface heat flux levels (W/ unit area, which is generally cm^2) thus become quite high in comparison to conventional heat rejection devices. The primary modes of heat transfer involve conduction through solid materials, convection which involves at least one moving gas or liquid, and "phase changes" which involves for example a liquid being converted to a vapor while absorbing energy from a hot body. Since the heat transfer coefficients occurring in the heat transport system will define the overall thermal resistance, it is desirable to have the highest possible values. In the similar manner heat transfer through any solid will induce a thermal resistance which is proportional to the material thickness and inversely proportional the material's thermal conductivity. These factors will emphasize the use of high thermal conductivity materials which posses reasonable mechanical strengths. While not directly associated with the heat transport system performance, the thermal expansion (contraction) coefficients of materials have to be carefully selected so as not to induce unacceptable stresses in the device or its supporting structure due to geometrical changes occurring over the wide temperature ranges experienced in cool down or heat up of the device.



While used extensively at lower heat fluxes, natural convection with a gas or liquid has limited cooling potential. Refrigerating the gas or liquid and then forcing it to flow over the device being cooled (forced convection) will dramatically improve cooling however the quantity of energy required to force the fluid through the device being cooled can become considerable. Also, the thermophysical properties of the fluid will come into play in defining the temperature difference which will exist between the cooled device and the fluid's temperature. At a specific flux level with a liquid coolant, the heat transfer properties of natural or forced convection will transition to those associated with a phase change such as boiling.

Figure IIe depicts the temperature difference, defined as device temperature minus fluid temperature for various heat transfer modes. While the fluids noted are more applicable to conventional and low temperature refrigeration systems, the general trends in temperature differences will hold for other fluids at lower temperatures. From this it is evident that for the higher heat flux levels of interest only convective or phase change heat transport systems will meet system requirements. Note, this does not mean that the conductive mode of heat transfer will not be used, only that it will be restricted to the use of highly conductive materials such as copper.

Figure II f shows the associated values of the heat transfer coefficients available via the heat exchange modes previously noted. Various low temperature refrigeration mixtures will have somewhat different values in both the forced convection and boiling regimes but will retain the trends noted. The boiling characteristic of liquid Nitrogen is somewhat similar to that of water. Helium represents the only unique gaseous heat transport medium that may find use in device cooling. Under the forced convection conditions, the heat transfer coefficients provided by Helium approach those of water without the temperature limits of water.

Infrared Components Corporation

January 13, 1999 - DABT63-97-C-00721



4.0 Refrigeration System Evaluation - Operating Temperatures, Performance, and Cost as a predecessor to down selection.

Temperature Range and Performance

Based on the information provided above it is possible to define the potential operating temperature regimes for various refrigeration systems. For the systems under consideration the question of cooling capacity is not likely to be an issue. All of the systems are capable of providing many 100's to 1000's of watts of cooling and are modular in nature. It is possible, and perhaps desirable, to use multiple refrigerator units to cool small groups of devices rather than a single large refrigerator cooling all the devices. This modularity will also have considerable market value since a single basic refrigeration system can be used in multiple applications. The performance (COP) of the cooling systems vary radically over the temperature range of interest which can lead to situations where the potential refrigeration system can provide the cooling capacity but the electrical input power requirements will be excessive. This is especially true throughout the Ultra-Low and Cryogenic temperature ranges.

An important computer system development issue, which cannot be fully defined in terms of temperature or performance, is the ability of the selected refrigeration system to be effectively employed over a wide range of temperatures and system configurations. Such a refrigeration system would only have to be "qualified" once with the initial cooled computer system. Operation at lower temperatures or in different computer configurations would not require development of a fundamentally different refrigeration cycle with the associated "qualification" process.

Figure IIg depicts the basic operating temperature ranges for the various refrigerator systems, while Figure IIh is an estimate of the cycle performance (COP) over the same temperature range.

From these figures it is evident that the refrigeration systems with the widest operating temperature range and best performance are the Mixed Gas / Expander versions of the Vapor Compression cycle and the Stirling cycle. The Stirling cycle systems are likely to be superior with regard to performance and physical size at temperatures below approximately -140°C . The two systems are performance and size competitive at temperatures up to approximately -50°C . The superiority of the Mixed Gas systems' intrinsic heat transports system (condensation at cold head - located remote from the compressor) will also play a significant role in defining the relative use of the two systems. A negative side of the Mixed Gas system is that at the lower temperatures, the refrigerant will likely contain combustible gas mixtures.

The simple conventional Vapor Compression cycle is clearly superior to the Mixed Gas systems at temperatures above approximately -80°C due to its reasonable performance, acceptable (at least in the US) refrigerant fluids, and lower cost. However, if temperatures below the -80°C range are desired, this system becomes quite complex and a transition to the Mixed Gas systems will be required.

Innovative Techniques for Cooling Computer Chips to Cryogenic Temperatures

Infrared Components Corporation
January 13, 1999 - DABT63-97-C-00721

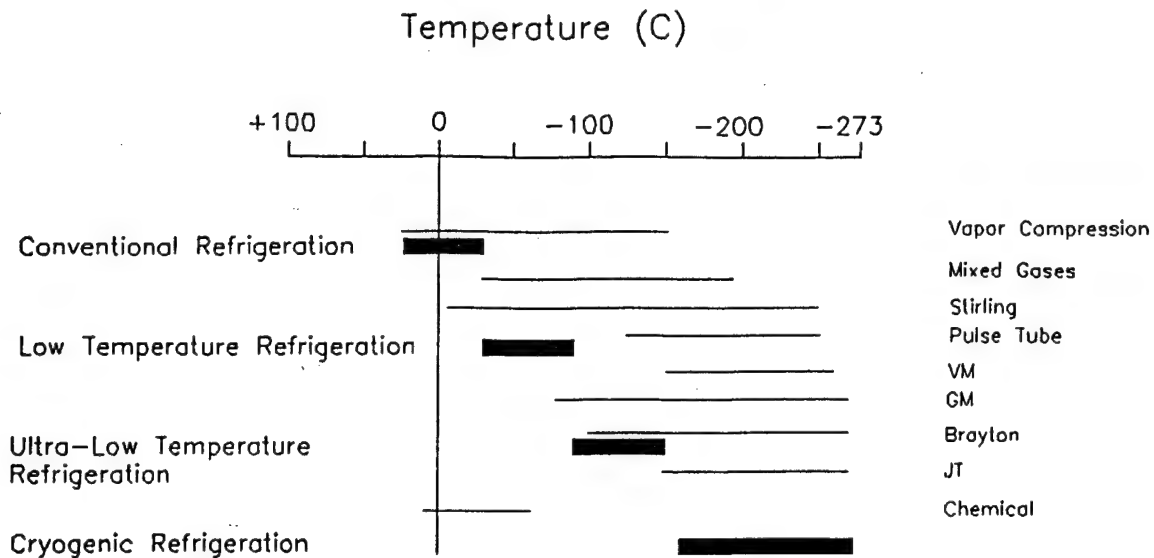


Figure IIg. Refrigeration System Operating Temperature Range




Figure IIh. Cooling Performance Vs Operating Temperature

The cooling of multiple CPUs can be carried out through the use of a single large refrigeration system or via the use of a number of smaller units which cool portions of the complete multi-CPU system. As previously noted, this latter approach has a number of potential system advantages, however in the case of units operating in the Ultra-Low and Cryogenic temperature regions, there exists a very strong economy of scale effect in performance of the refrigeration system as a function of capacity. For example, Figure J depicts the Percent of Carnot efficiency of various operating Stirling cycle coolers as a function of the cooling capacity at 77K. While the capacity range shown is much wider than expected in cooled computers, it is interesting to note that in the capacity range of 10 to perhaps 300 watts, the efficiency of the cooler nearly doubles. This falls well within the likely range of capacities for many of the systems under consideration. A similar effect will occur with Mixed Gas systems. GM systems tend to have even faster drop off in efficiency at smaller sizes (<5 watts) due to the nature of their compressor systems.

Refrigerator Cost Issues

Defining an accurate estimate of refrigeration cost effectiveness (cost/watt of cooling) represents a significant problem due to a number of factors. First and most critical is the fact that no fully-qualified, sub-ambient temperature computer cooling system has market experience. The DEC/KryoTech 767MHz system has just recently been introduced at a \$10k or 1/3 premium over the air-cooled baseline. Without sufficient experience, it is difficult to judge the additional costs which must be incorporated into the refrigeration system to meet computer system requirements. Also, as the operating temperatures are lowered, the number and types of refrigeration systems available fall dramatically. Many of these systems tend to be of a "laboratory" nature and as such their costs do not reflect the significant improvements that come with high sales volumes. Classic systems falling in this category are the Stirling cycle cooler and the Mixed-Gas systems which are marketed in small numbers for a variety of applications. The GM system has a market cost track record due to its wide use in cryopump applications in the electronics manufacturing industry, however even in this case, the issues of packaging the bulky compressor module within the computer system could lead to significant costs.

Low and Ultra-Low temperature systems based on the vapor compression cycles are not produced in large numbers, however they are fabricated from components which are utilized in the conventional temperature refrigeration systems. This provides relatively good refrigeration system cost effectiveness, however the cost of such units will not fall dramatically with increased sales volumes.

The material presented in Table II-2 is an attempt to estimate the refrigeration system cost effectiveness over various temperature ranges of interest, taking into account the issues noted above. The following assumptions have been utilized in this process:

1. The required cooling effect for an individual unit is in the range of 25 to 75 watts.
2. The useful temperature range is defined by the basic cycle characteristics.
3. Production volume is a minimum of 6,000 units per year.

4. The costs noted are manufacturing costs, including material and labor, but excluding profit or computer integration costs or auxiliary heat transport systems, if needed.

Because of the sensitivity of the results to the assumptions employed, and in some cases, the very small number of credible reference points available, the "spread" on the data is quite large, however it is useful for comparison between systems. While the specific cooling capacity selected for this evaluation is representative, it is interesting to note that the cost per watt is a function of capacity such that a 10 watt system is likely to cost \$100/watt while a well defined 250 watt system should cost on the order of \$35/watt.

Table II-2 Cost of cooling for various refrigeration cycles

Refrigeration Cycle	Temperature Range, °C	\$/Watt of cooling
Stirling	-200° to -80°	\$50 to \$100
Mixed Gas/Expander	-185° to -150°	\$60 to \$150
Gifford-MacMahon	-200° to -120°	\$40 to \$100
Multi-Stage Vapor Compression	-160° to -90°	\$15 to \$30
Single-Stage Vapor Compression	-70° to 0°	\$2 to \$10

5.0 Down-Selection at the 100W/-100°C operation point

Program specifications have been established for a net cooling capacity (CPU power dissipation and insulation losses) of 100 watts at an operating temperature nominally set to -100°C. This latter parameter will be varied over the range of -140°C to -85°C to reflect the "standard" vapor compression systems available in this temperature range. Over this temperature range the refrigeration systems also change from relatively simple 2-compressor cascade systems to considerably more complex mixed refrigerant auto cascade systems based on single compressors. These "standard" units are primarily employed in biomedical freezer applications and generally have cooling capacities greater than required in the current specifications, however these units do represent mass-produced products with high reliability requirements operating in the temperature range of interest. It is also critical to note these systems are made up of mass-produced refrigeration system compressors adapted for these temperatures.

The Stirling cycle refrigeration system does not possess a clear break point in its basic physical or operating characteristics as a function of cold head temperature. This attribute simplifies the thermodynamic analysis process. On the other hand, the likely requirement for a heat transport system coupling the Stirling cycle cold head to the cooled CPU assembly and the lack of an established manufacturing base makes the cost evaluation of this class of refrigeration system considerably less accurate than for vapor compression systems.

The primary criteria employed in the selection process for the proposed system focused on refrigeration systems cost, proven reliability, and ability to locate compressor "remote" from the cold head (evaporator) without experiencing excessive performance losses. The latter parameter was felt to be critical to the computer system designer due

to the packaging flexibility of this attribute.

Based on these factors, and over the temperature range of interest, the vapor compression system clearly is superior to the alternative Stirling cycle even if the cost and reliability of the two systems were equivalent. This is due to the ability to remotely locate the vapor compression system evaporator (cold head) at significant distance and orientations relative to the compressor system.

The fundamental CPU / cold head interface for both the vapor compression and Stirling systems are equivalent from a heat transfer viewpoint. This is due to the fact that the Stirling will have to employ some form of phase change heat transport system. The CPU "cold plate" or mounting surface thus is equivalent for either system.

5.1 Refrigeration system options:

Vapor Compression

Refrigerant Constraints - at the temperature range under investigation, there are a number of currently available refrigerants which are approved for world wide applications. Ever increasing restrictions Ozone layer effects will have an effect on some refrigerants in the future particularly for EU applications.

Compressor Lubrication - the -85°C systems have an excellent track record from the viewpoint of separating the compressor lubricant from the cold heat exchanger components. The mixed refrigerant systems have had problems with mixture "stratification" in time which can lead to migration of compressor lubricants to the cold portions of the system. This will lead to heat exchanger performance reductions and in time possibly to compressor failure.

Cascade vs Single Stage - for the temperature range of interest all single refrigerant systems will employ 2 compressors in a cascade configuration. The use of mixed refrigerant auto cascade systems does allow the use of a single compressor. Generally the mechanical complexity of this configuration restricts its use to the lower temperatures.

Fundamental Thermodynamic Performance - As in the case of all refrigeration cycles the performance (cooling capacity / power to drive cycle) falls with decreasing temperatures. Figure 11h depicts the relative power consumption for vapor compression systems over the temperature range noted. A -40°C system is used as a reference point since it would employ conventional "ice cream freezer" hardware. As can be seen the power consumption will be approximately twice that of the -40°C unit at the nominal -100°C operating temperature. This value is increasing rapidly as the temperature falls below -100°C.

Vibration - Since the cold head can be efficiently located well away from

the compressor, little or no vibration will be transmitted to the cold head. The cold head (evaporator) is intrinsically mechanically vibration free due to the absence of moving parts.

Cold Head (evaporator) heat transfer - the phase change processes occurring in the cold head (evaporator) represent the best possible heat transfer mode for an efficient, compact heat exchanger.

Cold Head location - The cold head can be located at significant distances from the compressor (5 to 10 ft) before a significant capacity reduction will be noted. This is a major technical advantage of this type of refrigeration system.

Technical Status - Fundamental technical base is well established for vapor compressions systems down to approximately -85°C based on two compressor cascade systems. Some questions exist on the more complex systems required to meet the lower temperatures under consideration.

Fundamental Cost Issues - In all vapor compression systems the cost of the system is dictated by the availability of commercial, high production volume refrigeration compressors. This represents a major advantage for this concept. It is highly unlikely that the number of compressors employed in any CPU cooling task will have any effect what so ever on the basic price of the compressors. This allows the initial cost to be relatively low in comparison to alternative refrigeration systems (Stirling) but the fundamental piece price will not fall dramatically with increased production volume of the CPU cooler system.

Depending on the compressor power requirements for the CPU cooling function it is also possible to be in the situation in which the refrigeration system must be made up of 2 compressors which have significantly more capability (and in turn physically much larger) than is actually required to meet the cooling requirement. This is due to the fact that nearly all "small" compressors are rated at 1/3, 1/2, and 1 Hp - there are no intermediate sizes. In addition smaller compressors can be significantly less efficient than larger units at the same operating temperatures. This is particularly true for the 1/3 Hp compressors which are likely to be the back bone of a vapor compression system operating at the capacities and temperatures of interest.

Technology Base - There is a problem with scaling conventional compressor to small capacity levels as may be required to meet package constraints. Lower temperature systems <-140°C represent a significant technology advancement.

Application Issues - Multi compressor systems will have high starting current requirements (relative to running requirements) especially if they are started hot. This likely will lead to a requirement for a separately

fused circuit for the compressor. Heat rejection from the refrigeration system becomes an important issue as the temperature is reduced due to the significant drop in system efficiency. This can lead to fan noise issues or minimum physical package size constraints.

Growth Potential

Higher or lower capacities (scaling) - higher capacities do not represent any problem for two compressor cascade systems (down to about -85°C) because of the ready availability of refrigeration compressors. Smaller sizes represent a problem as noted above.

Lower Operating Temperatures - Very limited, represents the major disadvantage of this cycle.

5.2 Stirling Cycle

Fundamental Thermodynamic Performance

Cold Head location - for all practical purposes, the cold head of the Stirling cycle refrigerator for this application will have to be located very near or be integral with the compressor system. At the capacities required in the proposed system, the losses in performance with separating the cold head cannot be tolerated.

Heat Transport Constraints (losses if utilized) - based on the evaluations performed to date it is likely that some form of heat transport system will be required to match the CPU requirements effectively to the Stirling cycle cold head operating characteristics. From the viewpoint of vibration isolation and the ability to separate the compressor (and Stirling cycle cold head) from the cooled electronics, the use of a heat transport actually represents a significant advantage to the Stirling cycle system. The use of a single, very small "chip" CPU may eliminate the need for the heat transport system since such a device can be effectively thermally coupled to the Stirling cold head. However the issues of vibration will come into play.

Practical, low cost passive heat transport systems for the Stirling cycle will have to incorporate some form of naturally circulating system (gravity and temperature driven). This will require that the condenser portion of the system (the Stirling cycle cold head) will have to be located above the cooled CPU so that the condensed liquid will flow down to the "hot" CPU, be vaporized and have the resulting gas flow up to the condenser due to buoyancy effects. Placing the relatively heavy Stirling system above the CPU may result in packaging problems.

Vibration - since major capacity reductions will occur if the cold head is placed at significant distances (>2 ft) from the compressor, some form of

vibration isolation system will be required.

Cold Head heat transfer - this represents a significant disadvantage for this cycle if a relatively large surface area (in terms of multiple CPU size) is to be cooled. Without the use of some form of heat transport system, the cooling will occur via conduction through solid materials to the helium working gas. This may be sufficient for a single very small "chip". These factors tend to drive the system to require some form of independent heat transport system with its added complexity and cost.

Growth Potential - The Stirling cycle has a major performance and size advantage over the vapor compression systems as temperatures fall to lower values. Even at the current target in the -100 to -140°C range, the Stirling cycle will be approximately 40 to 60 % more efficient than the vapor compression system discussed below. A potential advantage of the Stirling in the temperature range under consideration is its ability to be scaled to lower cooling capacities while retaining good overall efficiency. This could become an issue if the CPU cooling capacity requirements were to be reduced to the below the 40 to 60 W. At these capacities the number of available conventional compressors is limited and their physical size may cause problems in comparison to a Stirling unit.

Technology Base - The basic Stirling cycle technology as matured considerably in the last 7 years due to the advent of bearing / seal concepts which have removed a number of the life limiting characteristics of earlier devices. No fundamental technology barriers will eliminate the Stirling cycle refrigeration system from this application.

Fundamental Cost Issues - the Stirling cycle suffers from a poorly defined costs base due to the fact that essentially none have been produced in volume at the capacity levels of interest. Extensive cost evaluations by a number of reputable groups have shown that the Stirling cycle system can be very competitive with vapor compression systems operating at the lower end of the temperature range under consideration (<100°C) if sufficient number are produced. At higher temperatures (>-100°C), the two compressor cascade system is clearly superior from a cost viewpoint. Cost represents the major market barrier for the Stirling cycle unit to meet the system specifications.

5.3 Refrigeration System Characteristics

Figure III depicts a preliminary layout of a vapor compression system sized to meet the system specifications. Table II-3 lists a number of the fundamental characteristics of the system. The mechanical arrangement has the two compressors essentially stacked to provide a reasonable package foot print. The combined height of the two compressors was used to define the available heat exchanger height which was then combined with the unit width to define the heat exchanger surface area available. In general it was

found that this provided adequate area. It is important to note that the basic compressors and heat exchanger area are rated for a 90°F (32°C) ambient condition (air to the heat exchanger). A degradation of approximately 10% in cooling capacity occurs for each 10°F (5.5°C) temperature increase above this rating point. This will impact system packaging since it will be desirable for the refrigeration system heat exchanger to receive the coolest possible air. Passing the air through the computer first could lead to an important decrease in performance.

Table II-3 Reference Vapor Compression Refrigeration System

Parameter	Value or Range
Operating Temperature (C)	-85° to -100°C
Configuration	Two compressor cascade
Compressor size each (Hp)	1/5 to 1/4
Refrigerant	Likely R 502
Electrical Power Requirement (running) assuming 120 VAC / 60 Hz, watts	550 to 600 W
Starting current (90 F) amps	15 to 20 Amps
Heat Rejection	Air cooled, fan included
Physical Size (inches)	12 W x 16 L x 18 to 20 H
Total weight (#)	65 Lbs.




Figure III. Vapor Compression Refrigeration System Configuration

The preliminary sizing of the evaporator cold plate results in dimensions of approximately 4 x 4 inches, which must be capable of safely operating at charge pressures in excess of 750 psi. It was assumed that the cold plate is fabricated from two copper plates which are brazed together. The surfaces will be machined prior to brazing to form the flow passages for the refrigerant. The use of extremely fine surface features for the heat exchanger via the use of emerging micro machining techniques should be actively investigated. The active heat transfer area is expected to be approximately 75 % of the entire surface providing an "average" heat removal rate of 13,000 watts per square meter (13 w/cm²). Through proper tailoring of the flow passage geometry it will be possible to increase the heat removal rate by 2 to 3 times in very localized areas. This concept is highly dependent on the use of advanced machining concepts such as micro machining.

The physical structure of the cold plate will have to be quite robust so as to allow the CPU module to be intimately connected to the surface of the cold plate. At the heat fluxes noted above, any loss of good thermal contact could lead very rapidly to failure of the CPU due to over heating. This robust structure also has the advantage of providing

an intrinsic thermal spreading of the heating, resulting in somewhat more uniform temperatures.

The influence of the thermal insulation for the CPU package (and any connecting fluid lines) was not addressed in this preliminary analysis.

5.4 Cost Projections

It was essentially impossible to get accurate estimates of the cost for the refrigeration system portions of current medical storage units operating at the temperatures of interest. However, it was clear from the discussions that the majority of these costs were in the two compressors and the heat exchanger. The compressors were essentially stock units which are received by the medical refrigeration system assembler, attached (integrated with) the system heat exchanger, and charged with the desired refrigerant. The heat exchangers are custom made and produced in relatively low volumes. The proposed system for CPU cooling has approximately 1/3 to 1/2 the capacity of the higher quality medical storage refrigerators. This results in the need to use smaller, lower performing compressors to keep physical size of the system as small as possible. Proportionally these units are more expensive on a \$/watt basis than the larger systems (a factor of about 1.5 times).

Based on catalog (Grainger) costs for similar compressor equipment, the total cost of the compressors and associated heat exchanger / air cooling fan is in the range of \$500 to \$600. The relationship between these catalog costs and actual vendor costs is open to considerable interpretation and highly dependent on volume of units purchased. Nonetheless they do represent an effective reference point for overall system cost trade off studies.